DOCUMENT RESUME

ED 198 497 CS 005 898

AUTHOR Rumelhart, David E.
TITLE Understanding Understanding. Technical Report, December 1978 through December 1980.
INSTITUTION California Univ., San Diego. Center for Human Information Processing.
SPONS AGENCY National Science Foundation, Washington, D.C.; Office of Naval Research, Washington, D.C. Psychological Sciences Div.
REPORT NO CHIP-100; ONR-TR-8101
PUB DATE Jan 81
CONTRACT N00014-79-C-0323
GRANT NSF-BNS-76-15024
NOTE 39p.

EDRS PRICE MF01/PC02 Plus Postage.
DESCRIPTORS Adults; *Cognitive Processes; Critical Reading; *Hypothesis Testing; Psychological Studies; *Reading Comprehension; *Reading Processes; *Reading Research; Story Reading
IDENTIFIERS Schemata; *Schema Theory

ABSTRACT

The schema theory of comprehension as it relates to understanding while reading is reviewed in this report. Several techniques for the study of on-line measurement of comprehension are described that are based on a series of ten stories, each with two versions. The report notes that when subjects were asked to generate interpretations of stories while reading a story line-by-line, they hypothesized about the possible contents of the story and evaluated those hypotheses against the sentences as they read; if they found the new information confirmatory they maintained and further elaborated their hypotheses; and if they found the new information disconfirmatory they eliminated the hypotheses and constructed another consistent with the data. Also, the report notes that reading time measurements conducted on-line appeared to confirm the theory-generating view. The report concludes (1) that in order to understand the comprehension process a detailed description of the schemata that readers have available as well as an account of the conditions under which certain of these schemata are activated is needed, since there is a tremendous amount of knowledge brought to bear on even the simplest story comprehension task, and (2) that the process of comprehension is very much like the process of constructing a theory. (MKM)

****
* Reproductions supplied by EDRS are the best that can be made *
* from the original document. *
****

U S DEPARTMENT OF HEALTH,
EDUCATION & WELFARE
NATIONAL INSTITUTE OF
EDUCATION

THIS DOCUMENT HAS BEEN REPRODUCED EXACTLY AS RECEIVED FROM THE PERSON OR ORGANIZATION ORIGINATING IT. POINTS OF VIEW OR OPINIONS STATED DO NOT NECESSARILY REPRESENT OFFICIAL NATIONAL INSTITUTE OF EDUCATION POSITION OR POLICY

ED198497

Understanding Understanding

David E. Rumelhart

Program in Cognitive Science
Center for Human Information Processing
University of California, San Diego
La Jolla, California 92093

Research support was provided by the Personnel and Training Research Programs, Psychological Sciences Division, Office of Naval Research, under Contract Number N00014-79-C-0323, Contract Authority Identification Number, NR 157-437. Support was also provided by grant BNS-76-15024 from the National Science Foundation to D. E. Rumelhart. The views and conclusions contained in this document are those of the author and should not be interpreted as necessarily representing the official policies, either expressed or implied, of the Office of Naval Research. Approved for public release; distribution unlimited. Reproduction in whole or in part is permitted for any purpose of the United States Government.

CHIP 100
January 1981

Copyright © 1980 David E. Rumelhart

Preparation of this paper was supported by NSF grant BNS-76-15024 to D. E. Rumelhart, and by the Office of Naval Research under contract N00014-79-C-0323. Approved for public release; distribution unlimited.

CS 5 05 800

UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

REPORT DOCUMENTATION PAGE		READ INSTRUCTIONS BEFORE COMPLETING FORM
1. REPORT NUMBER 8101	2. GOVT ACCESSION NO.	3. RECIPIENT'S CATALOG NUMBER
4. TITLE (and Subtitle) Understanding Understanding		5. TYPE OF REPORT & PERIOD COVERED Technical Report 12/78 - 12/80
		6. PERFORMING ORG. REPORT NUMBER CHIP 100
7. AUTHOR(s) David E. Rumelhart		8. CONTRACT OR GRANT NUMBER(s) N00014-79-C-0323
9. PERFORMING ORGANIZATION NAME AND ADDRESS Center for Human Information Processing University of California San Diego La Jolla, California 92093		10. PROGRAM ELEMENT, PROJECT, TASK AREA & WORK UNIT NUMBERS NR 157-437
11. CONTROLLING OFFICE NAME AND ADDRESS Personnel and Training Research Programs Office of Naval Research (Code 458) Arlington, Virginia 22217		12. REPORT DATE January 1981 13. NUMBER OF PAGES 33
14. MONITORING AGENCY NAME & ADDRESS(if different from Controlling Office)		15. SECURITY CLASS. (of this report) Unclassified 15a. DECLASSIFICATION/DOWNGRADING SCHEDULE

16. DISTRIBUTION STATEMENT (of this Report)

Approved for public release; distribution unlimited.

17. DISTRIBUTION STATEMENT (of the abstract entered in Block 20, if different from Report)

18. SUPPLEMENTARY NOTES

19. KEY WORDS (Continue on reverse side if necessary and identify by block number)

Comprehension	Schemata
Understanding	Reading
Schema	Discourse

20. ABSTRACT (Continue on reverse side if necessary and identify by block number)

The schema theoretic view of comprehension is reviewed. Several techniques for the study of on-line measures of comprehension are developed. The techniques include a question probe technique and a reading time technique. Data were collected using each of these techniques and the results are compared with each other. The results are then interpreted in terms of a schema theoretical view of comprehension.

DD FORM 1 JAN 73 1473 EDITION OF 1 NOV 65 IS OBSOLETE
S/N 0102-LF-014-6601

Unclassified

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

Introduction

What is understanding? How do we make sense out of what we read or are told? I believe that over the past several years a substantial consensus has arisen in the field of Cognitive Science about the broad outlines of this process (cf. Fillmore, 1975; Minsky, 1975; Rumelhart, 1977; Rumelhart & Ortony, 1977; Schank & Abelson, 1977). In this paper I wish to sketch the basic features of those outlines and to show how this sketch can be given some reality by a careful analysis of the interpretations people actually make of stories and story fragments.

Consider the following brief fragment of a story:

> Mary heard the ice cream truck coming down the street. She remembered her birthday money and rushed into the house.

Upon hearing just these few words most readers already have a rather complete interpretation of the events in the story. Presumably Mary is a little girl who wants to buy some ice cream from the ice cream man and runs into the house to get her money. Of course, it doesn't say this in the story, there are other possibilities. Mary could be afraid that the ice cream man might steal her birthday money, etc. Still, most readers find the first interpretation most plausible and retain it unless later information contradicts it.

Consider, in contrast, the following story fragment.

> Mary heard the bus coming down the street. She remembered her birthday money and rushed into the house.

Upon hearing a fragment such as this, most people get a rather different notion of what the story might be about. The story fragment is less coherent. For most, Mary is older. Rather than the 4- to 8-year-old of the previous paragraph, Mary is now at least a teenager and possibly even an adult woman. Moreover, the quantity of money is somewhat greater. Almost surely the money is not needed to buy the passage on the bus itself--somehow bus fare is too mundane for birthday money.

Consider still another variation on this same story.

> Mary heard the ice cream truck coming down the street. She remembered her gun and rushed into the house.

Here we get a rather different interpretation again. Is Mary going to rob the ice cream man? Does she fear for her life? Note how the modification of a single word or phrase signals an entirely different interpretation. What sort of process could be accounting for such radical differences? Surely, it cannot be a process which takes word meanings and parlays them into sentence meanings and then those into text meanings.

The purpose of this paper is to explore the processes involved in these examples, to give a general account of these processes and to describe some experiments I have been doing in an attempt to understand them more fully.

To begin, let me lay out a general theoretical account of the comprehension process as I understand it and then turn to some data which, I believe, help explicate this process.

A Schema-Theoretic Model of Understanding

In my attempts to account for these phenomena I have found it useful to appeal to the notion of schemata. Before proceeding with a discussion of comprehension itself, it might be useful to explicate my notion of schemata.

A schema theory is basically a theory about knowledge. It is a theory about how knowledge is represented and about how that representation facilitates the use of the knowledge in particular ways. According to "schema theories" all knowledge is packaged into units. These units are the schemata. Embedded in these packets of knowledge is, in addition to the knowledge itself, information about how this knowledge is to be used.

A schema, then, is a data structure for representing the generic concepts stored in memory. There are schemata representing our knowledge about all concepts: those underlying objects, situations, events, sequences of events, actions, and sequences of actions. A schema contains, as part of its specification, the network of interrelations that is believed to normally hold among the constituents of the concept in question. A schema theory embodies a prototype theory of meaning. That is, inasmuch as a schema underlying a concept stored in memory corresponds to the meaning of that concept, meanings are encoded

in terms of the typical or normal situations or events which instantiate that concept.

Perhaps the central function of schemata is in the construction of an interpretation of an event, object or situation--in the process of comprehension. In all of this, it is useful to think of a schema as a kind of informal, private, unarticulated theory about the nature of the events, objects, or situations which we face. The total set of schemata we have available for interpreting our world in a sense constitutes our private theory of the nature of reality. The total set of schemata instantiated at a particular moment in time, constitutes our internal model of the situation we face at that moment in time, or, in the case of reading a text, a model of the situation depicted by the text.

Thus, just as the activity surrounding a theory is often focused on the evaluation of the theory and the comparison of the theory with observations we have made, so it is that the primary activity associated with a schema is the determination whether it gives an adequate account for some aspect of our current situation. Just as the determination that a particular theory accounts for some observed results involves the determinations of the parameters of the theory, so the determination that a particular configuration of schemata accounts for the data presently available to our senses requires the determination of the values of the variables of the schemata. If a promising schema fails to account for some aspect of a situation, one has the options of accepting the schema as adequate in spite of its flawed account or of rejecting the schema as inadequate and looking for another possibility. Therefore

the fundamental processes of comprehension are taken to be analogous with hypothesis testing, evaluation of goodness of fit, and parameter estimation. Thus, a reader of a text is presumably constantly evaluating hypotheses about the most plausible interpretation of the text. Readers are said to have understood the text when they are able to find a configuration of hypotheses (schemata) which offer a coherent account for the various aspects of the text. To the degree that a particular reader fails to find such a configuration, the text will appear disjointed and incomprehensible.

Schemata are like theories in another important respect. Theories, once they are moderately successful, become a source of predictions about unobserved events. Not all experiments are carried out. Not all possible observations are made. Instead, we use our theories to make inferences with some confidence about these unobserved events. So it is with schemata. We need not observe all aspects of a situation before we are willing to assume that some particular configuration of schemata offers a satisfactory account for that situation. Once we have accepted a configuration of schemata, the schemata themselves provide a richness which goes far beyond our observations. Upon deciding that we have seen an automobile, we assume that it has an engine, headlights, and all of the standard characteristics of an automobile. We do this without the slightest hesitation. We have complete confidence in our miniture theory. This allows our interpretations to far outstrip our observations. In fact, once we have determined that a particular schema accounts for some event we may not be able to determine which aspects of our beliefs are based on direct information and which are merely

consequences of our interpretation.

On Getting Some Evidence

I have been investigating story comprehension for several years and I have developed a story grammar (Rumelhart, 1975) which has proven rather useful in the analysis of story comprehension and recall. More recently (Rumelhart, 1977), I have recast that original work in the general framework described above and have developed a model capable of accounting for the kinds of summaries people give to very simple stories. Although this general approach to story understanding and story memory has proven rather popular, I have been dissatisfied with the work on two counts:

(1) Although much of the work (including my own) has focused on the process of story understanding, most of the experiments employed post-comprehension measures. Usually the measures have employed story recall and occasionally they have employed summarization. I have wished increasingly for truly 'on-line' measures of comprehension.

(2) The story grammar approach has tended to focus on rather abstract features of story comprehension. By its nature, the story schemata I and most others have studied offer a very general account of the structure readers see in stories. This generality is a plus in the sense that the schemata are very generally used, but they are a minus in the sense that they ignore the vast amount of other information which subjects can and do bring to bear in understanding stories.

During the last couple of years I have been attempting to develop some experimental techniques which could offer on-line information about subjects' comprehension processes. In the series of studies described in this paper, I set out to study this process of hypothesis generation and evaluation during the process of comprehension. Perhaps the simplest way to determine what people are thinking while they are

understanding is to ask them.

The basic experimental paradigm involved presenting subjects a series of stories a sentence at a time and, after each sentence, asking them WHO they thought the characters under discussion were, WHAT did they feel was going on in the story, WHY did the characters behave as they did, WHEN do they think the event described took place, and WHERE do they think the story is set. A series of 10 pairs of stories and/or story fragments were prepared. Most of the stories were based on initial segments of actual short stories written by well-known authors. The segments were edited slightly so that an alternate version of each story could be created through the modification of one or two words or phrases. The two story versions were designed, like the example story fragments at the beginning of this paper, so that the modification led to a rather different interpretation of the whole story. Each subject read one version of each one of the ten different stories. In order to assess the effects of the line at a time interpretation procedure on comprehension, some subjects were presented the stories two lines at a time, some four lines at a time, and still others were presented the whole story at one time.

There are two results which emerged immediately from this procedure:

(1) The process is very natural. Subjects report that is is very easy to describe the hypotheses that come to mind as they read. Unlike problem solving where the collecting of protocols seems to interfere with the process, our evidence indicates that, if anything, it actually <u>improves</u> comprehension.

(2) Subjects show a remarkable degree of agreement. With just three or four subjects the broad outlines of the sorts of

results generally obtained becomes clear.

Perhaps the best was to illustrate the procedure and the kinds of results obtained is by example. Consider the following sentence which is the first line from one of my stories.

> I was brought into a large white room and my eyes began to blink because the bright light hurt them.

Consider this sentence and what scene comes to mind. There was a good deal of agreement among my subjects. Almost without fail people believed that either this was an INTEROGATION situation in which the protagonist is being held prisoner, or it is a HOSPITAL scene in which the protagonist is a patient. It is also of some interest that when asked (after they had finished the story) why that had thought it was whatever they thought almost all reported that it was the bright lights or the large white room which had tipped them off. In point of fact, further experimentation seems to indicate that it was the "was brought" which was the key, putting protagonist in a passive situation. The large white room and bright lights simply further specify the basically passive situation aroused by the particular construction.

The OIL CRISIS story

As a second example, consider the following brief passage used in my experiment:

> Business had been slow since the oil crisis.
> Nobody seemed to want anything really elegant anymore.
> Suddenly the door opened and a well-dressed man entered the showroom floor.
> John put on his friendliest and most sincere expression and walked toward the man.

Although merely a fragment, my subjects generated a rather clear interpretation of this story. Apparently, John is a car salesman fallen on hard times. He probably sells rather large, elegant cars--most likely Cadillacs. Suddenly a good prospect enters the showroom where John works. John wants to make a sale. To do that he must make a good impression on the man. Therefore he tries to appear friendly and sincere. He also wants to talk to the man to deliver his sales pitch. Thus, he makes his way over to the man. Presumably, had the story continued John would have made the sales pitch and, if all went well, sold the man a car.

How, according to the theory described above, do people arrive at such an interpretation? Clearly, people do not arrive at it all at once. As the sentences are read, schemata are activated, evaluated, and refined or discarded. When people are asked to describe their various hypotheses as they read through the story, a remarkably consistent pattern of hypothesis generation and evaluation emerges. The first sentence is usually interpreted to mean that business is slow <u>because</u> of the oil crisis. Thus, people are led to see the story as about a business which is somehow dependent on oil as suffering. Frequent hypotheses involve either the selling of cars, or of gasoline. A few interpret the sentence as being about the economy in general. The second sentence, about people not wanting elegant things anymore, leads people with the gas station hypothesis into a quandary. Elegance just doesn't fit with gas stations. The gas station hypothesis is weakened, but not always rejected. On the other hand, people with hypotheses about the general economy or about cars have no trouble incorporating

this sentence into their emerging interpretation. In the former case they conclude it means that people don't buy luxury items and in the latter they assume it means that people don't buy large elegant cars--Cadillac's--much anymore. The third sentence clinches the car interpretation for nearly all readers. They are already looking for a business interpretation --that most probably means a SELLING interpretation-- and when a well dressed man enters the door he is immediately labeled as someone with MONEY-- a prospective BUYER. The phrase showroom floor clearly invalidates the gas station interpretation and strongly implicates automobiles which are often sold from a showroom. Moreover, the occurrence of a specific event doesn't fit at all well with the view that the passage is a general discussion of the state of the economy. Finally, with the introduction of John, we have an ideal candidate for the SELLER. John's actions are clearly those stereotypic of a salesperson. John wants to make a sale and his "putting on" is clearly an attempt on his part to "make a good impression." His movement toward the man fits nicely into this interpretation. If he is a salesman, he must make contact with the man and deliver the stereotypic "pitch."

Qualitatively, this account fits well with the general theoretical approach I have been outlining. The process of comprehension is very much like the process of constructing a theory, testing it against the data currently available, and as more data becomes available, specifying the theory further--i.e., refining the default values (as perhaps was the case when those holding the "car hypothesis" from the beginning encountered the sentence about nobody wanting anything elegant anymore). If the account becomes sufficiently strained, it is given up and a new

one constructed, or, alternatively, if a new theory presents itself which obviously gives a more cogent account, the old one can be dropped and the new one accepted.

But where do these theories come from? These theories are, of course, schemata. Presumably, through experience we have built up a vast repertoire of such schemata. We have schemata for car salesmen, the kinds of motives they have and the kinds of techniques they employ. We have schemata for automobiles, including how and where they are sold. We have built up schemata for the "oil crisis," what kinds of effects it has on what kinds of businesses. We have schemata about business people, the kinds of motives they have and the kinds of responses they make to these motives. The knowledge embedded in these schemata form the framework for our theories. It is some configuration of these schemata which ultimately form the basis for our understanding.

But how does a relevant schema suggest itself? Presumably, it is the "bottom-up" observation that a certain concept has been referenced that leads to the suggestion of the initial hypotheses. The notion that business was slow, suggests schemata about business and the economy. Since the slowness was dated from the occurrence of the oil crisis, it is a natural inference that the oil crisis was the cause of the slowness. Thus, a BUSINESS schema is activated. The particular TYPE of business is presumably a variable which must be filled. The information about the oil crisis suggests that it may be an oil - related business. Thus, readers are led to restrict the TYPE variable of the BUSINESS schema to oil-related businesses.

At this point, after the bottom-up activation of the high-level BUSINESS schema has occurred, this schema would generate a top-down activation of the various possible oil related businesses. Prime candidates for these are, of course, automobile related businesses. Of these, selling gasoline and automobiles are the two most salient possibilities.

When the second sentence is encountered, an attempt is made to fit it into the schemata currently considered most promising. As I discussed above, this information could serve to further restrict the TYPE variable in the automobile BUSINESS schema, but doesn't fit well with the gasoline business schema.

The BUSINESS schema presumably has a reference to the BUY or SELL schema. Once activated these schemata search for potential variable bindings. In the case of the automobile business, the MERCHANDISE variable is bound to an automobile. The second sentence suggests an elegant automobile. The reader has, when the third sentence is encountered not yet found a candidate for BUYER or SELLER. The sentence about a well-dressed man immediately suggests a potential BUYER. The phrase "showroom floor" offers additional bottom-up support for the automobile hypothesis. In fact, it is a strong enough clue itself that it can suggest automobile sales to a reader who currently considers an alternative schema more likely. We thus have a BUYER and some MERCHANDISE. The well-dressed quality of the BUYER is consistent with our view that the MERCHANDISE is elegant and therefore expensive--being well-dressed suggest MONEY. We need only a SELLER--i.e., an automobile salesman.

Readers probably already bring a relatively complete characterization of the "default value" for car salesman. We need but little additional information to generate a rather detailed description of goals and motives.

It is, in general, a difficult matter to analyze freeform responses of the sort obtained in this experiment. I have, however, devised a data representation scheme which allows the tracking of a subject's hypotheses through a story. The basic idea is illustrated in Figure 1. At any point in time a subject's hypothesis state can be characterized as a region in a multidimensional hypothesis space in which one dimension is time (or place in the story) and the other dimensions represents the subjects momentary beliefs about WHO the characters are, WHAT is going on in the story, WHERE the story is set, etc. Just two dimensions, WHERE and WHO, are illustrated in the figure. Each point in the space represents a possible hypothesis at some point in time. A particular subject's sequence of hypotheses can be represented as a path passing through the space. We can imagine that at particularly critical times during the reading of the story the path will turn sharply in several dimensions. At start, we might imagine that different subjects would occupy a fairly wide region of the space. By the end, all of the paths for the different subjects should have converged on one or two points in the space. Of course the dimensions other than the dimension of time are purely nominal and of course subjects often hold several hypotheses at once (i.e., they occupy not a point, but a region of the space), but nevertheless, this general representation proves useful in charting subjects' changing hypotheses.

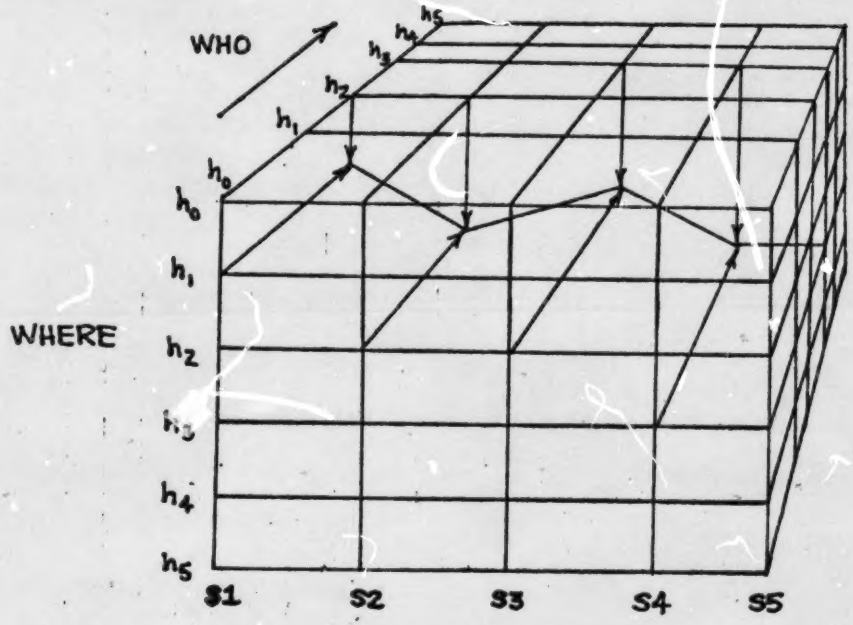

Figure 1. A representation of a subject's shifting hypotheses while reading a story. One dimension represents the sequence of sentences in the story. The other two dimensions represents a subject's hypotheses with respect to WHO the characters are and WHERE the action is taking place. The vector passing through the space represents a possible sequence of hypotheses.

I will illustrate the general form of analysis by looking at some of the results from the "Oil Crisis" story. In order to analyze the data, the responses for each question were categorized and for each subject it was recorded which of the responses was given. For example, there were five different categories of answers to WHERE the story took place. The five categories were:

(1) Indefinite: when subjects said they had no clear idea.

(2) Gas station: when subjects believed that the action was occurring at a gas station.

(3) Showroom: when subjects believed that the action took place in a automobile showroom.

(4) Luxury store: when subjects believed that the action took place in a luxury store such as a jewelry store or a fancy furniture or clothing store.

(5) Nation: when subjects believed that the story was a general statement about the national economy.

Figure 2 illustrates the patterns of responses observed from the ten subjects who read this version of the story. Each line on the graph represents a pattern of responses. The number on the lines represents the number of subjects showing that pattern. We can see that five subjects had no clear idea where the events were taking place after the first sentence. One subject thought from the start that it was in an automobile showroom. Four subjects thought, after the first sentence, that the story was taking place in a gas station. We can see that after the second sentence four people moved to the view that it was an automobile showroom, three thought it was in a luxury store, two were still indefinite and one thought it was a general discussion of a national economy.

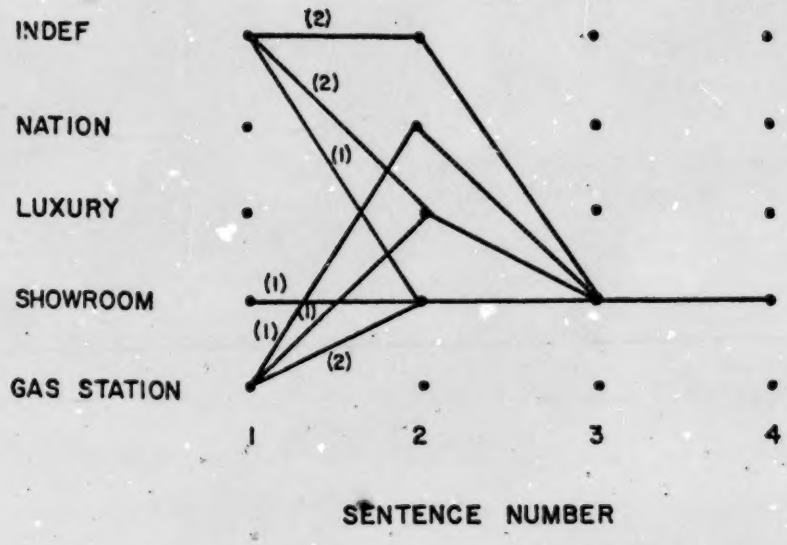

Figure 2. Set of paths through the hypothesis space for the question of WHERE the "Oil Crisis" story was taking place.

There is not space here to illustrate the whole pattern of results for this story, rather, I turn now to a discussion of a second story which shows a more dramatic pattern of results.

The DEAR LITTLE THING story

Consider now the following story used in my experiment:

(1) Dear little thing.
(2) It was nice to feel it again.
(3) She had taken it out of its box that afternoon, given it a good brush and rubbed life back into its dim little eyes.
(4) Little rogue! Yes, she really felt that way about it.
(5) She put it on.
(6) Little rogue, biting its tail just by her left ear.
(7) When she breathed something gentle seemed to move on her bosom.
(8) The day was cool and she was glad she had decided on her little fur.

The results for this story are particularly interesting. As people read the story they form clear impressions of certain aspects of the story, but none of them consider the possibility that the story might be about a fur until the fifth line of the story and for some, this is not clear until the last line of the story. From the beginning, however, many readers have an impression that the speaker in the story is a woman. Of the twenty people to read the first line of the story, seven mentioned that they thought the it was a woman speaking. In none of my other stories did people spontaneously assign a sex to the speaker after only reading the first sentence. Apparently a number of the readers interpret the pattern of speech here to be typically feminine. This is illustrative of the subtlety of the kinds of clues readers pick up on and that authors count on.

Perhaps the most interesting response was that which subjects made to the WHAT questions. Here we get the clearest picture of their overall assessment of what the story is about. There were six categories of responses given by our subjects. These were:

(1) CLOTHING: they thought that the woman was talking about a hat or some jewelry.

(2) FUR: They thought the woman was talking about a fur.

(3) LETTER: They thought someone was writing a letter.

(4) PET: They thought the story was about a pet.

(5) STIMULATION: They thought the story was about sexual stimulation.

(6) TOY: They thought the story was about a stuffed animal or doll.

Figure 3 shows the pattern of hypotheses held by the ten people who read this version of the story. After the opening line, "Dear little thing," people were about evenly split between the possibility that it was about a pet or letter writing. The second line, "It was nice to feel it again," discouraged all but one of the letter-writing hypotheses. Some of these decided that it was a "toy" or stuffed animal that the story was about. Others assumed it was about sexual stimulation or had no clear idea. The third line moved almost everyone who didn't think it was a "pet" to the view it was a "toy." The fourth line offered no new information and people held on to their previous hypotheses. The fifth line, "She put it on" was difficult to assimilate with any of the hypotheses and, as is evident from the figure, nearly everyone switched to the view that it was either a piece of clothing or jewelry or to the view that it was a fur piece. The seventh line strengthened the FUR hypothesis and the eighth line clinched it for everyone.

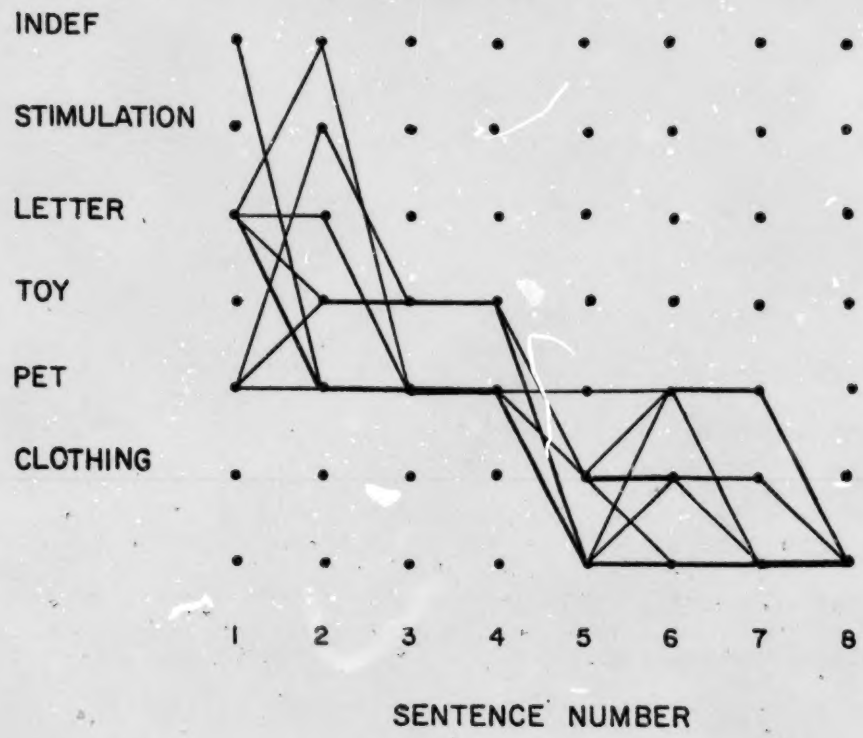

Figure 3. Set of paths through the hypothesis space for the question of WHAT in the "Dear Little Thing " story.

The figure clearly shows the critical nature of the fifth sentence. We can see subjects, on the basis of such "bottom up" information as the use of the word "dear," determine that it might be a letter or a diminutive reference to a pet. Then, once finding a satisfactory hypothesis, maintaining and refining it until disconfirming information is made available. Then, when disconfirmation occurs, searching out a new workable hypothesis.

Clearly, in this case, my subjects are behaving according to the hypothesis evaluation mode that I have suggested. But, is this the normal way of processing? Doesn't the procedure force them to respond in this way? These are serious questions. Indeed, I do believe that there is an effect of the procedure. However, I believe that it is better categorized as making subjects read more carefully than at modifying the basic procedure. One bit of evidence for this view is that over all of the stories, subjects who interpreted the stories a line at a time more often agreed with each other (and with the experimenters) about the interpretation of the story than subjects who gave an interpretation only after having read the whole story. In addition, a second experiment was carried out to try to get an alternative measure of "on-line" processing. In this experiment the subjects were not asked to make any interpretations of the story. Rather, they were presented the story one word at a time and asked to press a button after they read each word to get the next word. The time to read each word was recorded. We can then compare different versions of the same story, one in which we know from the "interpretation" experiments requires a rather dramatic shift in hypotheses, and another which requires no such shift or a shift at a

different place. The "Dear Little Thing" story offers an ideal example. The alternative version of this story differed in three words. Sentence 5 was "She put it down" rather than "She put it on." Sentence 6 ended "by her left ankle" rather than "by her left ear," and sentence 8 ended "take her pet along" rather than "take her fur along." Thus, for one version, the FUR version, subjects probably had to shift hypotheses after line 5. For the other version, the PET version, subjects probably already had the correct hypothesis by the line 5. Thus, the two stories were identical for the first 49 words and differed in only three of the final 38 words.

Since we know from the interpretation experiment that a good deal of re-evaluation occurs in the FUR form of the story after line 5 and that a large number of subjects have the PET hypothesis well before line 5, it is reasonable that people would read the last 38 words of the story more slowly in the FUR version. Table 1 shows the average reading time per word for the first 49 and last 38 words for the two versions of the story. The expected difference is apparent in the table. The average reading time for the first half of the story is about the same for the two groups. Those with the FUR version were about 200 msec slower over the last half of the story. Unfortunately, the magnitudes in the table are probably somewhat misleading. There is an average difference of some 20 milliseconds between the groups for the first half of the story. In fact, this average is a mixture of some early slow responses and some later fast responses for the PET group. A better estimate for the difference between the two groups base-line reading-speed is 125 milliseconds per word. Thus, the apparent 200 milliseconds per word

Table 1

Mean Reading Time per Word in Milliseconds

Story Version	First 49 Words	Last 38 Words
FUR	886	1011
PET	864	801

difference evident in the table is probably actually closer to a 75 millisecond difference per word. Nevertheless, even with this conservative estimate of the difference between the base reading rate of the two groups, the 75 msec per word figure over the 38 words amounts to an average difference of almost 3 seconds longer for the FUR group. Thus, inspite of some difficulties with the data here, it would appear that we have been able to see, in slower reading times, the same hypothesis re-evaluation our subjects in the interpretation experiment told us about.

A somewhat closer look at the data appear to confirm this conclusion. Much of this effect is already evident on the reading of the last word of line 5. Figure 4 shows the reading times for each word in the line. The most obvious characteristic of these curves is the increased reaction time for the last word of the sentence for both versions of the story. This upswing on the last word of a sentence is normal in all of experiments of this sort. It appears to represent some sort of "consolidation" phase of the reading process. More important to the present discussion, however, is the difference in response time between those subjects who heard the word "on" and those who heard the word "down" as the last word of the sentence. Upon hearing the word <u>on</u> the PET or TOY hypotheses are disconfirmed and subjects are forced to begin to re-evaluate their hypotheses. This re-evaluation apparently takes time. Indeed, as Figure 5 indicates, many of the subjects are apparently still formulating more hypotheses through the following sentence. Notice, for example, the time required by the subjects with the FUR version as compared with those on the PET version for the word <u>tail</u>. Presumably those with the PET version have already hit upon the pet hypothesis and thus

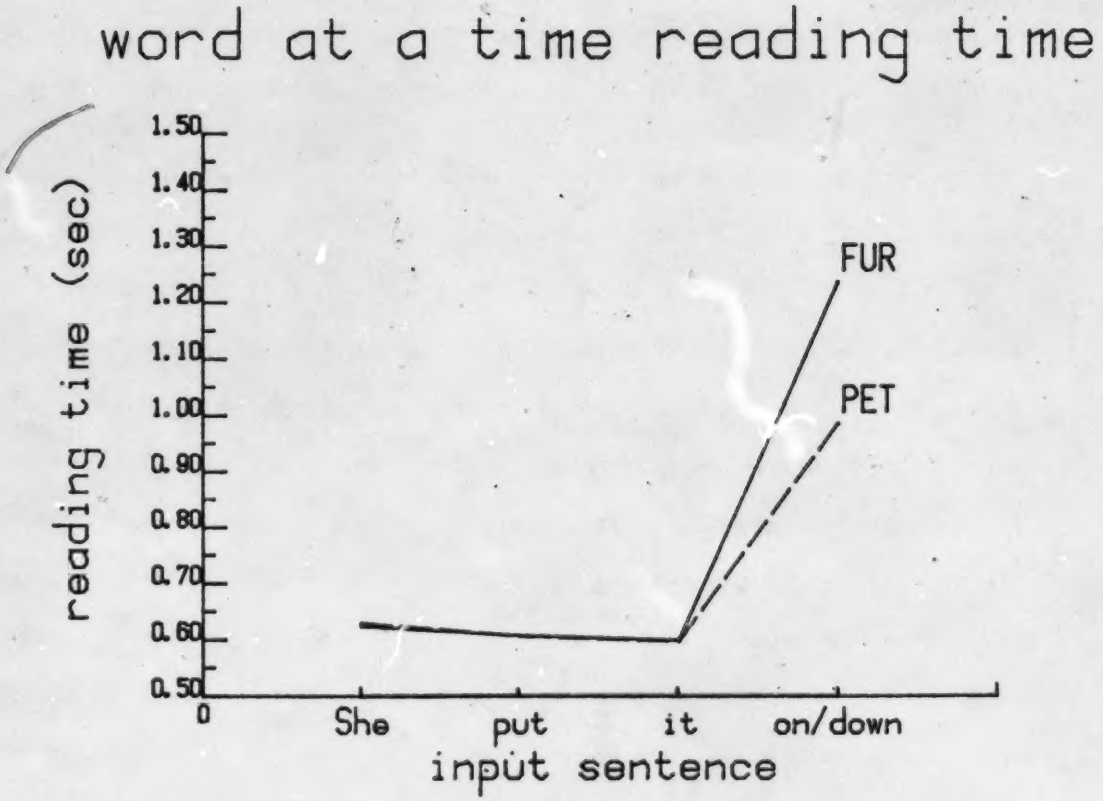

Figure 4. Adjusted Word by word reading times for the two versions of line 5 of the "Dear Little Thing" story. Due to overall differences in the reading rates of the two groups, the times for the FUR group were adjusted downward by subtracting 125 msec for each point. This value was chosen so that the two groups showed about the same level of performance over the three words before the two stories diverge.

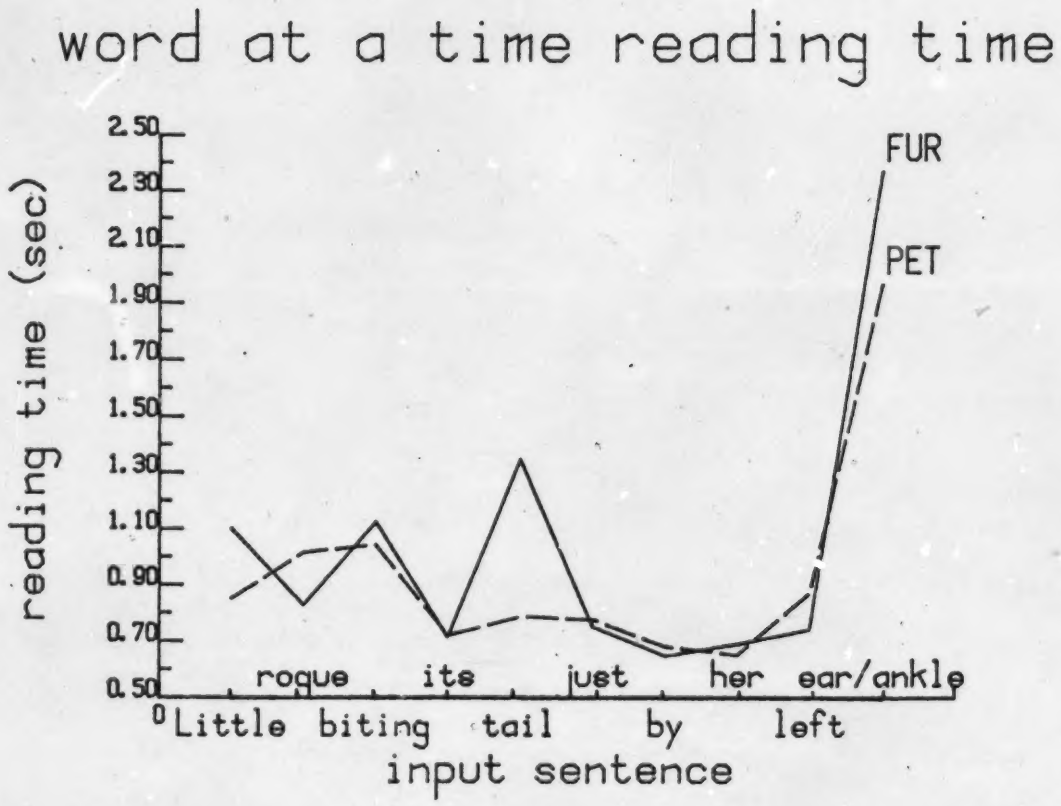

Figure 5. Adjusted word by word reading times for the two versions of line 6 of the "Dear Little Thing" story. Again the times for the PET data represent raw data while the FUR curve has been adjusted downward by 125 msec to adjust for overall differences in reading time between the two groups which was evident before the stories diverged.

the word tail fits nicely into their existing interpretation. Many of those in the FUR version have probably opted for the hypothesis that the story is about a piece of clothing or jewelry and thus are not able to integrate "tail" into their existing interpretation-- similarly, for the last word of the sentence. The subjects with the PET version have little or no trouble with the pet being near the woman's ankle. The FUR subjects find it difficult to reconcile something with a tail being near the woman's ear.

Overall, in spite of the unfortunate baseline differences between the two groups, the reading time results does appear to confirm the view that a very different method of gaining access to on-line processing leads to a generally congruent pattern of results.

General Comments

I have tried, in this section, to present a flavor of the results I have been collecting in the context of story comprehension. Due to limitations of space I have been unable to present a complete analysis of all of my data. Nevertheless, these examples should serve to illustrate the major points. When asked to generate interpretations of stories while reading through a story line-by-line, subjects generate hypotheses about the possible contents of the story and evaluate them against the sentences as they read them. If they find the new information confirmatory they maintain and further elaborate their hypotheses. If they find the new information disconfirmatory they eliminate the hypothesis and construct another consistent with the input data. This process seems to involve both "top-down" and "bottom-up" processes. Certain words and

phrases appear to "suggest" from the bottom-up certain frameworks of interpretation--such as the INTEROGATION framework in the first example. Once a particular interpretation has received a moderate degree of support, it can come to guide the processing and interpretation of future inputs. Subjects find it relatively easy and natural to go rather substantially beyond the specifics of what the input sentences actually say. Their interpretations contain material about aspects of the situation which are totally unaddressed in the input text.

To what degree is this a natural process, as the schema theory sketched in the first section suggests, and to what degree does the procedure force the subjects to behave as the theory suggests they would? This is a difficult question for this approach. It is extraordinarily difficult to get data which bear on this issue. Three approaches have been tried:

(1) I have collected word by word reading times for subjects not instructed to generate interpretations and have looked for correlations between points in the story where we believe subjects to be evaluating new hypotheses and those where we observe elevated response times. By and large, as the examples presented above illustrate, these two measures correlate.

(2) I have collected interpretations of subjects after they read the whole story and compared them with those of subjects who read the stories a line at a time. The results showed that subjects who interpreted a line at a time nearly always generated the same interpretations as those who gave us an after-the-fact interpretation. The only discernable difference was that those who gave an interpretation only at the end showed somewhat more variability in their interpretations. It appears that this results from more careless reading on the part of the subjects offering an interpretation only at the end.

(3) I have asked a few subjects for retrospective analyses of the processes they went through while reading the stories immediately after reading the stories. Although such subjects mention fewer hypothesis changes than those giving on-line

interpretations, the overall structure of their reports seem to parallel those of the on-line subjects.

None of these methods is really totally convincing in and of itself. Nevertheless, the combination of the fact that the response times seem to follow the hypothesis interpretations, the fact that the interpretation paradigm doesn't seem to affect the final interpretations subjects generate and the fact that in informal observations subjects' retrospective reports seem very similar to the line-at-a-time results points strongly to the view that the general pattern of hypothesis generation observed in our experiments is present in normal reading.

On Understanding and Misunderstanding

Before concluding, it is useful to consider the application of this general theory to the notion of misunderstanding. On the present account, understanding is the process of finding a configuration of schemata which offers an adequate account of a passage or situation. The analysis given above illustrates how such a process is supposed to operate. Clues from the story suggest possible (instantiations of schemata) which are then evaluated against the successive sentences of the story until finally a consistent interpretation is discovered. Sometimes, a reader fails to correctly understand a passage. There are at least three reasons implicit in schema-theory as to why this might occur.

(1) The reader may not have the appropriate schemata. In this case he/she simply cannot understand the concept being communicated.

(2) The reader may have the appropriate schemata, but the clues provided by the author may be insufficient to suggest them. Here again the reader will not understand the text, but, with

appropriate additional clues may come to understand it.

(3) The reader may find a consistent interpretation of the text, but may not find the one intended by the author. In this case, the reader will "understand" the text, but will misunderstand the author.

There are many examples of these three phenomena in the literature. Perhaps the most interesting set of studies along these lines were carried out by Bransford and Johnson (1973). They studied the comprehension of texts in which subjects lacked the appropriate schemata, ones in which the schemata were potentially available, but there were not sufficient clues to suggest the correct ones as well as ones in which subjects were led to choose a "wrong" interpretation. Consider the following paragraph used in one of their studies.

> The procedure is actually quite simple. First you arrange things into different groups. Of course, one pile may be sufficient depending on how much there is to do. If you have to go somewhere else due to lack of facilities that is the next step, otherwise you are pretty well set. It is important not to overdo things. That is, it is better to do too few things at once than too many. In the short run this may not seem important but complications can easily arise. A mistake can be expensive as well. At first the whole procedure will seem complicated. Soon, however, it will become just another facet of life. It is difficult to foresee any end to the necessity for this task in the immediate future, but then one can never tell. After the procedure is completed one arranges the materials into different groups again. Then they can be put into their appropriate places. Eventually they will be used once more and the whole cycle will then have to be repeated. However, that is part of life. [p. 400]

Most readers find this passage, as written, extremely difficult to understand. However, once they are told that it is about washing clothes, they are able to bring their clothes-washing schema to the fore and make sense out of the story. The difficulty with this passage is thus not that readers don't have the appropriate schemata, rather, it

stems from the fact that the clues in the story never seem to suggest the appropriate schemata in the first place. The "bottom-up" information is inadequate to initiate the comprehension process appropriately. Once the appropriate schemata are suggested, most people have no trouble understanding the text.

Although most readers simply find the passage incomprehensible, some find alternative schemata to account for it and thus render it comprehensible. Perhaps the most interesting interpretation I have collected was from a Washington bureaucrat who had no difficulty with the passage. He was able to interpret the passage as a clear description of his job. He was, in fact, surprised to find that it was supposed to be about "washing clothes" and not about "pushing papers." Here then, we have an example of the third kind of comprehension failure, "understanding the story," but "misunderstanding the author."

Conclusion

At this point, it might be useful to put this comprehension theory in the context of a theory of communication. I find it useful, in this regard, to think of the general view of comprehension put forth here as suggesting that the problem facing a comprehender is analogous to the problem that a detective faces when trying to solve a crime. In both cases there are a set of clues. The listener's (or reader's) job is to find a consistent interpretation of these clues. In so doing, the listener uses his or her own prior experiences and knowledge of the speaker to create a most plausible possibility. Just as the meaning of a particular clue that a detective might find cannot be determined

except in relation to the way it fits into the whole situation, so to the meaning of a particular word, phrase, or even sentence cannot be interpreted except in relation to the way it fits into the whole of the story. Similarly, from the speaker's point of view. The speaker's (or writer's) problem is to leave a trail of clues which, in the opinion of the speaker, will lead the reader to make the inferences that the speaker wishes to communicate. Thus, the speakers must use their knowledge of the listener or, at least of the cultural expectations of the listener, to create the set of clues which most reliably and economically leads the listener to the desired hypotheses.

Thus, the authors of short stories need not spell out every detail. Instead, they provide the reader with subtle clues which they can expect the reader will pick up on. Thus, in the example of the INTEROGATION scene the author, by subtle use of the passive and the mention of bright lights and a white room has generated in the reader a full-blown image of an entire INTEROGATION scene. The remainder of the story can then play off of these subtle clues and needn't waste time or words setting the scene. Similarly, in the "Dear Little Thing" story the author has, in a single phrase, suggested to many a woman speaking. I suspect that these stories are not at all unusual. I suspect that in general all of the inferences we wish to communicate can <u>never</u> be "spelled out" and that we must <u>always</u> depend on our ability to draw forth the appropriate schemata in the listener through a large variety of clues.

Finally, let me comment on the direction I wish to push the sort of work I have discussed here. I have, for the past several years. been

attempting to create a computer simulation system capable of comprehending language according to the kinds of principles just described. I have taken as an empirical goal the attempt to create a program capable of mimicking the experimental results from the interpretation experiments. Obviously, a detailed account of the comprehension process requires a detailed description of the schemata readers have available as well as an account of the conditions under which certain of these schemata are activated. There is a startling amount of knowledge brought to bear on even the simplest story comprehension task. Nevertheless, I believe that data of the sort I have described above will provide a useful data base against which to evaluate models of comprehension.

References

Fillmore, C. J. An alternative to checklist theories of meaning. Proceedings of the First Annual Meeting of the Berkeley Linguistics Society, 1975, 1, 123-131.

Minsky, M. A framework for representing knowledge. In P. H. Winston (Ed.), The psychology of computer vision. New York: McGraw-Hill, 1975.

Rumelhart, D. E. Notes on a schema for stories. In D. G. Bobrow & A. M. Collins (Eds.), Representation and understanding: Studies in cognitive science. New York: Academic Press, 1975.

Rumelhart, D. E. Understanding and summarizing brief stories. In D. La Berge & S. J. Samuels (Eds.), Basic processes in reading: Perception and comprehension. Hillsdale, N. J.: Erlbaum Associates, 1977.

Rumelhart, D. E., & Ortony, A. The representation of knowledge in memory. In R. C. Anderson, R. J. Spiro, & W. E. Montague (Eds.), Schooling and the acquisition of knowledge. Hillsdale, N. J.: Erlbaum Associates, 1977.

Schank, R., & Abelson, R. Scripts, plans, goals and understanding: An inquiry into human knowledge structures. Hillsdale, N. J.: Erlbaum Associates, 1977.

CHIP Technical Report List

1. David M. Green and William J. McGill. On the equivalence of detection probabilities and well known statistical quantities. October, 1969.
2. Donald A. Norman. Comments on the information structure of memory. October, 1969.
3. Norman H. Anderson. Functional measurement and psychophysical judgment. October, 1969.
4. James C. Shanteau. An additive decision-making model for sequential estimation and inference judgments. October, 1969.
5. Norman H. Anderson. Averaging model applied to the size-weight illusion. October, 1969.
6. Norman H. Anderson and James C. Shanteau. Information integration in risky decision making. November, 1969.
7. George Mandler, Richard H. Meltzer, Zena Pearlstone. The structure of recognition. Effects of list tags and of acoustic and semantic confusion. November, 1969.
8. Dominic W. Massaro. Perceptual processes and forgetting in memory tasks. January, 1970.
9. Daniel Graboi. Searching for targets: The effects of specific practice. February, 1970.
10. James H. Patterson and David M. Green. Discrimination of transient signals having identical energy spectra. February, 1970.
11. Donald A. Norman. Remembrance of things past. June, 1970.
12. Norman H. Anderson. Integration theory and attitude change. August, 1970.
13. A.D. Baddeley and J.R. Ecob. Reaction time and short-term memory: A trace strength alternative to the high-speed exhaustive scanning hypothesis. November, 1970.
14. A.D. Baddeley. Retrieval rules and semantic coding in short-term memory. December, 1970.
15. Roy D. Patterson. Residue pitch as a function of the number and relative phase of the component sinusoids. March, 1971.
16. George Mandler and Marilyn A. Borges. Effects of list differentiation, category membership and prior recall on recognition. May, 1971.
17. David E. Rumelhart, Peter H. Lindsay, and Donald A. Norman. A process model for long-term memory. May, 1971.
18. David E. Rumelhart and Adele A. Abrahamson. Toward a theory of analogical reasoning. July, 1971.
19. Martin F. Kaplan. How response dispositions integrate with stimulus information. August, 1971.
20. Martin F. Kaplan and Norman H. Anderson. Comparison of information integration and reinforcement models for interpersonal attraction. August, 1971.
21. David M. Green and R. Duncan Luce. Speed-accuracy tradeoff in auditory detection. September, 1971.
22. David E. Rumelhart. A multicomponent theory of confusion among briefly exposed alphabetic characters. November, 1971.
23. Norman H. Anderson and Arthur J. Farkas. New light on order effects in attitude change. March, 1972.
24. Norman H. Anderson. Information integration theory: A brief survey. April, 1972.
25. Donald A. Norman. Memory, knowledge, and the answering of questions. May, 1972.
26. David J. Weiss. Averaging: An empirical validity criterion for magnitude estimation.
Norman H. Anderson. Cross-task validation of functional measurement. June, 1972.
27. David E. Rumelhart and Patricia Siple. The process of recognizing tachistoscopically presented words. August, 1972.
28. Ebbe B. Ebbesen and Richard Bowers. The effects of proportion of risky to conservative arguments in a group discussion on risky shift. September, 1972.
29. Ebbe B. Ebbesen and Michael Haney. Flirting with death: Variables affecting risk taking on our nation's highways. September, 1972.
30. Norman H. Anderson. Algebraic models in perception. November, 1972.
31. Norman H. Anderson. Cognitive algebra: Information integration applied to social attribution. December, 1972.
32. Jean M. Mandler and Nancy L. Stein. Recall recognition of pictures by children as a function of organization and of distractor similarity. January, 1973.
33. David E. Rumelhart and Donald A. Norman. Active semantic networks as a model of human memory.
Marc Eisenstadt and Yaakov Kareev. Towards a model of human game playing. June, 1973.
34. George Mandler. Memory storage and retrieval: Some limits on the reach of attention and consciousness. July, 1973.
35. Kent L. Norman. A method of maximum likelihood estimation for stimulus integration theory. August, 1973.
36. Yaakov Kareev. A model of human game playing. August, 1973.
37. Donald A. Norman. Cognitive organization and learning. August, 1973.
38. The Center for Human Information Processing: A Five Year Report — 1968-73.

39. Larry D. Rosen and J. Edward Russo. Binary processing in multi-alternative choice. October, 1973.

40. Samuel Himmelfarb and Norman H. Anderson. Integration theory analysis of opinion attribution. December, 1973.

41. George Mandler. Consciousness: Respectable, useful, and probably necessary. March, 1974.

42. Norman H. Anderson. The problem of change-of-meaning. June, 1974.

43. Norman H. Anderson. Methods for studying information integration. June, 1974.

44. Norman H. Anderson. Basic experiments in person perception. June, 1974.

45. Norman H. Anderson. Algebraic models for information integration. June, 1974.

46. Ebbe B. Ebbesen and Vladimir J. Konečni. Cognitive algebra in legal decision making. September, 1974.

47. Norman H. Anderson. Equity judgments as information integration.
Arthur J. Farkas and Norman H. Anderson. Input summation and equity summation in multi-cue equity judgments. December, 1974.

48. George Mandler and Arthur Graesser II. Dimensional analysis and the locus of organization. January, 1975.

49. James L. McClelland. Preliminary letter identification in the perception of words and nonwords. April, 1975.

50. Donald A. Norman and Daniel G. Bobrow. On the role of active memory processes in perception and cognition. May, 1975.

51. J. Edward Russo. The value of unit price information. An information processing analysis of point-of-purchase decisions. June, 1975.

52. Elissa L. Newport. Motherese: The speech of mothers to young children. August, 1975.

53. Norman H. Anderson and Cheryl C. Graesser. An information integration analysis of attitude change in group discussion. September, 1975.

54. Lynn A. Cooper. Demonstration of a mental analog of an external rotation.
Lynn A. Cooper and Peter Podgorny. Mental transformations and visual comparison processes: Effects of complexity and similarity. October, 1975.

55. David E. Rumelhart and Andrew Ortony. The representation of knowledge in memory. January, 1976.

56. David E. Rumelhart. Toward an interactive model of reading. March, 1976.

57. Jean M. Mandler, Nancy S. Johnson, and Marsha DeForest. A structural analysis of stories and their recall: From "Once upon a time" to "Happily ever after". March, 1976.

58. David E. Rumelhart. Understanding and summarizing brief stories. April, 1976.

59. Lynn A. Cooper and Roger N. Shepard. Transformations on representations of objects in space. April, 1976.

60. George Mandler. Some attempts to study the rotation and reversal of integrated motor patterns. May, 1976.

61. Norman H. Anderson. Problems in using analysis of variance in balance theory. June, 1976.

62. Norman H. Anderson. Social perception and cognition. July, 1976.

63. David E. Rumelhart and Donald A. Norman. Accretion, tuning and restructuring: Three modes of learning. August, 1976.

64. George Mandler. Memory research reconsidered: A critical view of traditional methods and distinctions. September, 1976.

65. Norman H. Anderson and Michael D. Klitzner. Measurement of motivation.
Michael D. Klitzner and Norman H. Anderson. Motivation x expectancy x value: A functional measurement approach. November, 1976.

66. Vladimir J. Konečni. Some social, emotional, and cognitive determinants of aesthetic preference for melodies in complexity. December, 1976.

67. Hugh Mehan, Courtney B. Cazden, LaDonna Coles, Sue Fisher, Nick Maroules. The social organization of classroom lessons. December, 1976.

67a. Hugh Mehan, Courtney B. Cazden, LaDonna Coles, Sue Fisher, Nick Maroules. Appendices to the social organization of classroom lessons. December, 1976.

68. Norman H. Anderson. Integration theory applied to cognitive responses and attitudes. December, 1976.

69. Norman H. Anderson and Diane O. Cuneo. The height + width rule in children's judgments of quantity. June, 1977.
Norman H. Anderson and Clifford H. Butzin. Children's judgments of equity. June, 1977.

70. Donald R. Gentner and Donald A. Norman. The FLOW tutor: Schemas for tutoring. June, 1977.

71. George Mandler. Organization and repetition: An extension of organizational principles with special reference to rote learning. May, 1977.

72. Manuel Leon. Coordination of intent and consequence information in children's moral judgements. August, 1977.

73. Ted Supalla and Elissa L. Newport. How many seats in a chair? The derivation of nouns and verbs in American Sign Language. November, 1977.

74. Donald A. Norman and Daniel G. Bobrow. Descriptions: A basis for memory acquisition and retrieval. November, 1977.

75. Michael D. Williams. The process of retrieval from very long term memory. September, 1978.

76. Jean M. Mandler. Categorical and schematic organization in memory. October, 1978.

77. James L. McClelland. On time relations of mental processes: A framework for analyzing processes in cascade. October, 1978.

78. Jean M. Mandler and Marsha DeForest. Developmental invariance in story recall. November, 1978.
Jean M. Mandler, Sylvia Scribner, Michael Cole, and Marsha DeForest. Cross-cultural invariance in story recall. November, 1978.

79. David E. Rumelhart. Schemata: The building blocks of cognition. December, 1978.

80. Nancy S. Johnson and Jean M. Mandler. A tale of two structures: Underlying and surface forms in stories. January, 1979.

81. David E. Rumelhart. Analogical processes and procedural representations. February, 1979.

82. Ross A. Bott. A study of complex learning: Theory and methodologies. March, 1979.

83. Laboratory of Comparative Human Cognition. Toward a unified approach to problems of culture and cognition. May, 1979.

84. George Mandler and Lawrence W. Barsalou. Steady state memory: What does the one-shot experiment assess? May, 1979.

85. Norman H. Anderson. Introduction to cognitive algebra. June, 1979.

86. Edited by Michael Cole, Edwin Hutchins, James Levin and Naomi Miyake. Naturalistic problem solving and microcomputers. Report of a Conference. June, 1979.

87. Donald A. Norman. Twelve issues for cognitive science. October, 1979.

88. Donald A. Norman. Slips of the mind and an outline for a theory of action. November, 1979.

89. The Center for Human Information Processing: A Description and a Five-Year Report (1974-1979). November, 1979.

90. Michael Cole and Peg Griffin. Cultural amplifiers reconsidered. December, 1979.

91. James L. McClelland and David E. Rumelhart. An interactive activation model of the effect of context in perception. Part I. April, 1980.

92. James L. McClelland and J.K. O'Regan. The role of expectations in the use of peripheral visual information in reading. February, 1980.

93. Edwin Hutchins. Conceptual structures of Caroline Island navigation. May, 1980.

94. Friedrich Wilkening and Norman H. Anderson. Comparison of two rule assessment methodologies for studying cognitive development. June, 1980.

95. David E. Rumelhart and James L. McClelland. An interactive activation model of the effect of context in perception. Part II. August, 1980.

96. Jean M. Mandler. Structural invariants in development. September, 1980.

97. David E. Rumelhart and Donald A. Norman. Analogical processes in learning. October, 1980.

98. James A. Levin and Yaakov Kareev. Personal computers and education: The challenge to schools. November, 1980.

99. Donald A. Norman and Tim Shallice. Attention to action: Willed and automatic control of behavior. December, 1980.

100. David E. Rumelhart. Understanding understanding. January, 1981.

Note: Requests for CHIP reports should be addressed to the author. Reports are also available through the Library Loan Service of the University of California, San Diego, La Jolla, California 92093.